

Main page: [Cisco Unified MeetingPlace, Release 8.0](#)

Up one level: [Planning Your Deployment](#)

Use the information in this section to determine which type of Application Server failover you want for your Cisco Unified MeetingPlace system.

- [How Application Server Failover Works](#)
- [Requirements for Application Server Failover](#)
- [Application Server Failover Deployments](#)
- [Failover for the Express Media Server](#)

Contents

- [1 How Application Server Failover Works](#)
- [2 Requirements for Application Server Failover](#)
- [3 Application Server Failover Deployments](#)
 - ◆ [3.1 About Configuring Your Application Server for Intra-Site Failover](#)
 - ◇ [3.1.1 Figure: Intra-Site Failover](#)
 - ◇ [3.1.2 Restrictions and Requirements for Intra-Site Failover](#)
 - ◆ [3.2 About Configuring Your Application Server for Inter-Site Failover with Reservationless Single Number Access \(RSNA\) Deployments](#)
 - ◇ [3.2.1 Figure: Inter-Site Failover](#)
 - ◇ [3.2.2 Related Topics](#)
- [4 Failover for the Express Media Server](#)

How Application Server Failover Works

MeetingPlace Application Server redundancy is available with any deployment model: Webex Scheduling or MeetingPlace Scheduling. This is also available with either Media server type: Hardware media server or Express Media Server.

For Application Server failover, you deploy two Application Servers—an active and a standby—as part of one logical site. Each Application Server in a site is called a node. Only one Application Server is active at any time. The active Application Server runs all processes and subsystems, while the standby Application Server only runs the database and Apache Tomcat subsystems.

Each Application Server has an Informix database that contains information about users, groups, and meetings. The system uses the Informix Enterprise Replication (ER) feature to allow changes in the database of one Application Server to be automatically carried over to the database of the other Application Server.

If the active Application Server fails, you can activate the standby Application Server and place the previously active Application Server in standby mode. You activate the standby Application Server either by using SSH remotely or directly on the server. This requires manual intervention by someone with administration privileges.

If your deployment includes a one-to-one redundancy of Hardware Media Servers, associate one Hardware

Media Server with the active Application Server and the redundant Hardware Media Server with the standby Application Server. If the primary and standby Application Servers share a common Hardware Media Server, add the Hardware Media Server to both Applications Servers.

Requirements for Application Server Failover

- The standby Application Server must have its own set of licenses. The licenses used for the active Application Server do not work on the standby Application Server.
- The time must be synchronized between both sets of Application Servers and between the Application Server and the web server.
- The active and the standby Application Servers can reside in different physical locations if the following two conditions are met:
 - ◆ They must both be configured on the same VLAN.
 - ◆ There is a dedicated failover Hardware Media Server.
- Failover requires a minimum bandwidth of 4 mbps.
- There is no minimum requirement for latency, which affects the delay in getting the data replicated to the other server.
- The network sends a heartbeat between the two replicating databases every 60 seconds. If there is no response, the network pings the database from which there was no response. If the database still does not respond, the system buffers all data until the heartbeat detects that the database is back. The network then sends the updated information to the database that was not previously responding.
- If an Application Server will not be available for a long time, for example, in case of a hardware failure, turn off replication between the Application Servers and reestablish replication after both Application Servers become available. This prevents the replication queue from becoming full.

Note: If the replication queue for an Application Server becomes full, the system shuts down. The system sends alarms when the replication queue becomes 75 percent full and again when it becomes 95 percent full, but there is no automatic mechanism to switch off replication.

Application Server Failover Deployments

You can configure your system for Application Server failover in one of two ways:

- Using intra-site failover, which does not allow changes to be made to both databases at the same time. See the [About Configuring Your Application Server for Intra-Site Failover](#).
- Using inter-site failover, which does allow you to make changes to both databases at the same time. See the [About Configuring Your Application Server for Inter-Site Failover with Reservationless Single Number Access \(RSNA\) Deployments](#).

The inter-site failover deployment only applies to RSNA (Reservationless Single Number Access) deployments. For more information about RSNA, see the [Reservationless Single Number Access \(RSNA\) Deployment](#).

Cisco Unified MeetingPlace does not support combined intra-site and inter-site failover.

About Configuring Your Application Server for Intra-Site Failover


In this deployment, there is one site with two Application Servers, or nodes, that share data related to users, groups, and meetings, but only one Application Server can be active at any time. There is a maximum of two

Application Servers allowed per site.

The system replicates the data between the two databases. The system does not allow simultaneous changes to database tables that contain users, groups, and meeting information.

If there is a network disconnection between the two Application Servers and certain records are modified in the databases for both Application Servers, the system uses the update timestamp to resolve the conflict. When the connection is restored, the latest update is propagated to both the servers.

Figure: Intra-Site Failover



Restrictions and Requirements for Intra-Site Failover

- The time must be synchronized between the two Application Servers. This is required to resolve conflicts when the same piece of data is modified simultaneously in both Application Servers.
- For failover deployments, the primary and standby Application Servers communicate using TCP port 2008 for database replication purposes. The system must allow communication on TCP port 2008 between the servers.
- Allocate the following hostnames and IP addresses:
 - ◆ Node 1 eth0 and Node 2 eth0-Use the same hostname and IP address for the eth0 network interface on both Node 1 and Node 2. Example: meetings.example.com, 10.0.0.1
 - ◆ Node 1 eth0:1-Use a unique hostname and IP address for this virtual network interface. Example: meetings1.example.com, 10.0.0.2
 - ◆ Node 2 eth0:1-Use a unique hostname and IP address for this virtual network interface. Example: meetings2.example.com, 10.0.0.3
- Make sure that the Domain Name System (DNS) server is configured for all three hostname-IP address pairs.

About Configuring Your Application Server for Inter-Site Failover with Reservationless Single Number Access (RSNA) Deployments

Note: This deployment only applies to RSNA (Reservationless Single Number Access) deployments. For more information about RSNA, see the [Reservationless Single Number Access \(RSNA\) Deployment](#).


This deployment allows two simultaneously active and inter-connected Application Servers, or nodes, at two different sites. This configuration is useful when the Application Servers need to have the same set of user

Cisco_Unified_MeetingPlace_Release_8.0_--_About_Failover_Options_for_the_Application_Server profiles (from an LDAP server) and user groups.

One of the Application Servers gets the user data from the LDAP server through Cisco Unified Communications Manager and then the Application Servers are synchronized through database replication.

Other than the user profile and user group synchronization, the Application Servers are independent and are associated with different web servers. A user connecting to an Application Server will see different configuration data (for example, meetings) except for the user profiles and user groups. This deployment allows users to conduct reservationless meetings on the other server, if one server fails.

Figure: Inter-Site Failover



Related Topics

For information on configuring the Application Server for failover, see the online help in the administrator interface or the *Configuration Guide for Cisco Unified MeetingPlace Release 8.0* at http://www.cisco.com/en/US/products/sw/ps5664/ps5669/products_installation_and_configuration_guides_list.html.

Failover for the Express Media Server

The Express Media Server is a software-based media server that resides on the Application Server. Therefore, failover for the Express Media Server is the same as for the Application Server.